LEVEL

NOSC

D D C

MAY 2 1979

RECEIVED

Technical Report 354

EVALUATION OF THE TWO-DIMENSIONAL FOURIER TRANSFORM APPLIED TO ACOUSTIC IMAGE PATTERN MEASUREMENTS

RS French

30 November 1978

Final Report: October 1977 – June 1978

Prepared for
Naval Electronic Systems Command

**NAVAL OCEAN SYSTEMS CENTER
SAN DIEGO, CALIFORNIA 92152**

79 05 01 007

NOSC

NAVAL OCEAN SYSTEMS CENTER, SAN DIEGO, CA 92152

AN ACTIVITY OF THE NAVAL MATERIAL COMMAND

RR GAVAZZI, CAPT, USN
Commander

HL BLOOD
Technical Director

ADMINISTRATIVE INFORMATION

ACKNOWLEDGEMENTS

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NOSC/TR-354	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* EVALUATION OF THE TWO-DIMENSIONAL FOURIER TRANSFORM APPLIED TO ACOUSTIC IMAGE PATTERN MEASUREMENTS		5. TYPE OF REPORT & PERIOD COVERED Final Report. October 1977 – June 1978
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Robert S. French		8. CONTRACT OR GRANT NUMBER(s) F11101
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Ocean Systems Center San Diego, CA 92152		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 62711N XF11 101 100
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Electronic Systems Command Washington, D. C.		12. REPORT DATE 30 November 1978
		13. NUMBER OF PAGES 34
14. MONITORING AGENCY NAME & ADDRESS*(if different from Controlling Office)*		15. SECURITY CLASS. *(of this report)* UNCLASSIFIED
	36p.	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

The objective of this experimental study was to evaluate the use of the two-dimensional Fourier transform as an aid in extracting spatial frequency components in acoustic data images by means of direct operator observation.

More specifically, the study addressed the question of the relative gain in precision of measurement which might be realized using the two-dimensional transform as compared to making these measurements directly from the original image. Of particular interest was the amount of gain at the weaker signal strengths typical of the class of processed acoustic data under investigation.

next page

393 159

20. Continued

As the signal power in the image decreases relative to the noise power, the probability increases that energy concentrations resulting from random noise patterning will appear in the transformed image. It is possible, therefore, that the human operator working with the original image alone may be able to separate signal-related patterning from spurious noise patterning as well as can be achieved by using the transform. The experiment reported herein was designed to provide a definitive test of this hypothesis.

Results showed that at high SNR levels the transform domain image representation was clearly superior to the original image as a mode for systematic measurement and ranking of periodic spatial features in the image. As the SNR level was reduced the accuracy of measurement in the transform domain fell off more rapidly than the original image domain, resulting in a convergence of the performance curves under the two modes. With the SNR level reduced so that the signal-related patterning was barely perceptible, measurement of periodic spatial features in the transform domain was no better than in the original image domain.

SUMMARY

PROBLEM

Two-dimensional transform processes are a promising new tool for acoustic image analysis since they provide decomposition of the image in terms of its two-dimensional spatial frequency components. These components are readily identified in the transform domain by noting the locations of major energy concentrations. Previous studies have shown that the representation of an acoustic image in the form of its two-dimensional Fourier transform enhances the identification and measurement of periodic spatial features in the image. However, none of the previous work has made a quantitative comparison of the accuracy of measurements in the transform domain with those made from the original image alone. In this report the effectiveness of pattern feature measurement is addressed by comparing operator performance using the transform image versus the original image over a range of SNR levels, including much lower SNR levels than previously studied.

RESULTS

At high SNR levels, the transform domain image representation is clearly superior to the original image as a mode for systematic measurement and ranking of periodic spatial features in the image. As the SNR level is reduced the accuracy of measurement in the transform domain falls off more rapidly than in the original image domain, resulting in a convergence of the performance curves under the two modes. With the SNR level reduced so that the signal-related patterning is barely perceptible, measurement of periodic spatial features in the transform domain is no better than in the original image domain.

A supplementary analysis was performed to identify factors that might account for the superior performance in the transform mode at the higher SNR levels. The magnitude or prominence of the spatial feature was found to be such a factor. At each SNR level, performance scores for the three most prominent spatial features were compared with the scores for three less prominent features. At the highest SNR level in the original image mode, the measurement accuracy for the less prominent features was significantly lower than for the three most prominent features. A comparable difference in accuracy was not found when the measurement accuracies in the transform domain were compared, i.e., at the highest SNR level the less prominent spatial features were, relative to the most prominent features, more difficult to identify and measure in the original image domain than in the transform domain. At the two lower SNR levels no comparable differences between the two modes were observed. The measurement accuracy for the three less prominent periodic features was equally poor in both the original image and transform domain.

1

CONTENTS

INTRODUCTION

BACKGROUND

The technology area of digital image processing has matured considerably during recent years, particularly in applications dealing with digitized natural images. However, many image processing concepts can be directly applied to other classes of digital images, e.g., in the area of acoustic signal displays, ambiguity surfaces, frequency-azimuth data arrays, and lofargrams.[1] This report provides an experimental assessment of the application of digital image processing to a class of acoustic imagery characterized by strong patterning.

The basic hypothesis which motivated this investigation is that the tools of image processing can be applied to acoustic signal space imagery to enhance the perceptibility and measurement of periodic spatial features. Since the two-dimensional Fourier transform process decomposes an image in terms of its two-dimensional spatial frequency components, the transform provides a measure of the relative energy for all spatial frequencies in the original image. From this measure, spatial features in the original image can be individually identified and analyzed based on the location and magnitude of the spatial frequency energy concentrations.

OBJECTIVE

The objective of this experimental study was to evaluate the use of the two-dimensional Fourier transform as an aid in extracting spatial frequency components in acoustic data images by means of direct operator observation.

More specifically, the study addressed the question of the relative gain in precision of measurement which might be realized using the two-dimensional transform as compared to making these measurements directly from the original image. Of particular interest was the amount of gain at the weaker signal strengths typical of the class of processed acoustic data under investigation.

As the signal power in the image decreases relative to the noise power, the probability increases that energy concentrations resulting from random noise patterning will appear in the transformed image. It is possible, therefore, that the human operator working with the original image alone may be able to separate signal-related patterning from spurious noise patterning as well as can be achieved by using the transform. The experiment reported herein was designed to provide a definitive test of this hypothesis.

---

[1]Naval Ocean Systems Center, NOSC TR 337, "Two-Dimensional Transform Processing for Acoustic Signal Space Imagery," by J. A. Roese, August 1978.

APPROACH

DATA DISPLAY FORMATS

The acoustic images used to evaluate the transform process were displayed in color on the COMTAL display system (appendix A). Operator performance was compared using two data display formats: the original image alone and the transform image representation in combination with the original image. An example of the combination format for one of the test patterns is shown in figure 1. In this figure the format presents data in four quadrants. Note that the image transform itself is displayed in a rectangular coordinate system forming four subquadrants.

The image transform in figure 1 provides a logarithmically scaled and normalized representation of the image energy magnitude coded in color. Phase information is not displayed in this format.

The location of concentration of energy or magnitude peaks in the transform permits direct measurements of spatial frequencies (number of horizontal and vertical axis crossings) of any periodic components in the original image. These spatial frequencies may be determined from the transform by noting the X and Y coordinates for any given peak. Note that the data displayed in the transform are symmetric with respect to the origin of the horizontal and vertical axes. Thus, all data needed for interpreting the transform energy distribution are available in the two quadrants to the right of the center axis.

The lower left quadrant of figure 1 contains data parameters which are readout automatically by the computer. Parameters displayed include the horizontal and vertical spatial frequency values, the slope of the periodic component, and the relative magnitude of any point within the images directly under the cursor. The use of these parameters will be described in more detail in a later section.

EXPERIMENTAL VARIABLES

The primary independent variable in the experiment was the mode in which the data were displayed together with the interactive measurement aids associated with each mode. The test patterns consisted of processed data from three different sources. Calibrated amounts of Gaussian noise at two levels were added to each of the original data sets, forming test sets of three images each in which the signal-to-noise ratio (SNR) varied from +5.8 to −9.3 dB. The three different test patterns plus the variation in SNR for each set were parametric variables, i.e., variables held constant while the primary variables were changed to determine their effects. The dependent performance variable was a score based on the accuracy and completeness of the subject's measurement of the parameters of each of the six strongest periodic components relative to a set of reference values estimating the "true" coordinates of these components. These reference values were derived from measurements in the transform domain for each image with no added noise.

INTERACTIVE MEASUREMENT AIDS

Measurement aids were available for both the original image and the transform modes of presentation. When the original image alone was viewed, the subject was able to obtain exact measurement of the spatial frequency and slope by constructing an overlay, or grid of lines, with any desired spacing and slope. An example of such an overlay corresponding to the predominant periodic component for one pattern is in figure 2. Note the readout of spatial frequency components and slope. In this case the magnitude value shown is a readout of the point directly beneath the cursor at the time the overlay parameters were entered into the computer. An overlay such as this is constructed by the subject by entering three points into the computer using the trackball and cursor. The first two points are used to define a line coincident with a major spatial feature while the third point defines the desired spacing between lines. The task for the subject in the original image mode is to determine predominant spatial frequency components (up to six) and measure their spatial frequency coordinates by constructing overlays with lines which coincide with each of the perceived spatial patterns. The parameters obtained for each overlay are recorded manually on a prepared data sheet.

When both the transform and the original image are present, the subject has additional interactive aids available. If the subject wishes to determine a periodic component from the image transform, the cursor is placed directly over a major energy concentration or peak and an overlay and readout are requested. The computer displays an overlay on the original image corresponding to the spatial frequency coordinates of the selected peak and also draws an "X" over the selected peak in the transform for future reference. In addition, the data parameters for the coordinates of the peak are displayed. An example of an overlay constructed in such a fashion is in figure 3. Because the transform image does not display the phase of the periodic component, the resultant overlay is not necessarily coincident with the pattern feature, although it has the correct spatial frequency and slope. An additional interactive feature, using the cursor, allows the subject to reposition the overlay. A line of the overlay can be made coincident with any desired peak or ridge in the original image. Thus the subject is able to establish the correspondence between major peaks in the transform and corresponding spatial features in the original image. The subject can also construct any desired overlay in the original image with a corresponding "X" overlay in the transform. This coupling of the transform to the original image allows the subject to verify that the parameters of the constructed overlay correspond to a high magnitude energy peak. This aids the subject in selecting the six strongest spatial frequency components as required by the test instructions.

A graphic representation of two periodic spatial features with grid overlays as they might appear in the image and transform domains is in figure 4. Note that the spatial frequency is determined from the image domain by the number of horizontal and vertical axis crossings of each grid, counting from the lower right corner of the image. In the transform domain the spatial frequency of each feature is represented by the location of the corresponding energy peak, scaled in X,Y coordinates that correspond to the number of horizontal and vertical axis crossings. For grids inclined to the left of vertical axis, the sign for the Y coordinate is negative and the energy peak will always appear in the upper of the two quadrants for the transform domain. In a similar manner, grids inclined to the right of the vertical axis have a positive slope with spatial frequency coordinates in the lower quadrants of the transform domain. Because the transform representation is symmetric with respect to the origin, only two of the four quadrants are shown.

8

CALIBRATED TEST DATA SETS

Three SNR levels were selected for test from a more complete set of 11 calibrated images that had been used in a previous study.[2] The levels selected were the highest SNR, an intermediate value, and the level at the threshold of pattern perception. Construction of the calibrated test data sets is described in detail in reference 2. Some of these details are repeated here. The original images were derived from a processed segment of real data with a high SNR. For this case the SNR can be calculated by using the maximum amplitude of the image. With this case as a reference, known amounts of noise power were added to the original image to produce a readily determined change in SNR. The same additive noise wave form was used for each SNR to control any effect caused by noise patterning. The SNR is defined for these experiments as $10 \log_{10}$ (signal power to total noise power) in a 1-Hz band at the input to the signal-processing algorithm. One test data set with SNR levels of +5.8, −5.1, and −9.3 dB is in figure 5, which illustrates the effects of noise on the perceptibility of the periodic components that are readily apparent in the original image. Note that the image color coding was not normalized to match the dynamic amplitude range of the image as the SNR was decreased. Since normalization of the output color coding improves pattern perceptibility in the original image, lack of normalization is a factor which should bias the results in favor of the transform mode.[2]

EXPERIMENTAL DESIGN

The design for this experiment is usually called a "subject by treatments" design, since each subject is tested under all conditions. In effect, each subject acts as his/her own control. Such a design is the most efficient test available for a relatively small, diverse group of subjects. The design for the experiment consisted of all combinations of the following variables: 3 (image patterns) $\times$ 2 (presentation modes) $\times$ 3 (SNR levels) = 18 trials. The order of presentation during the 18 trials was the same for each subject. This order is most easily described in terms of the rapidity with which each variable is changed. The most rapidly changing variable was pattern, then the two modes, and finally the SNR. The order of mode presentation was original image alone followed by the combination of transform and original image. The three SNR levels were presented in order from lowest to highest.

TEST PROCEDURES

The general procedure for the experiment required the subjects to determine and record the spatial frequency coordinates of the six strongest periodic components in each image for each trial. The subjects determined the coordinates for each spatial feature with the aid of overlays. The vertical and horizontal coordinates and slope value were recorded manually on a special data sheet. The test procedure was straightforward. Prior to starting the test the subject read the written instructions (appendix B). There were no practice trials. Each subject had served in a previous study using the same test data patterns. They were also familiar with the use of the cursor and the data terminal keyboard. For the first

---

[2]Naval Ocean Systems Center, NOSC TR 207, "Grey-Scale Versus Color Coding of Acoustic Data Images," by Robert S. French, March 1978.

few trials the test conductor remained to answer procedural questions and to determine that the subject was operating the equipment properly. To minimize carryover from one set of measurements to the next, the subjects were not allowed to compare measurements for a given pattern under the two modes or SNR conditions. Once a given data sheet was completed, the subject was not allowed to change any entry or to refer to the sheet again during the test.

SCORING

SCORING ALGORITHM

The subject's recorded measurements, consisting of groups of six spatial frequency coordinates for each of the 18 trials, were scored in terms of their accuracy and completeness relative to the coordinates of a set of reference values. Each pair of coordinates for a given spatial feature was assigned a score based on the absolute "miss" distance relative to the "true" coordinates for the periodic spatial feature in question. This general concept of scoring is illustrated in figure 6, using the same pattern and coordinates used for feature 1 in figure 4. It should be emphasized that the difference between the "true" coordinates and the subject's corresponding overlay for this particular spatial feature has been greatly exaggerated for illustrative purposes and is not typical of the test results.

The subject's performance score for a given trial was based on the average miss distance ($\underline{d}$) for the six X,Y coordinate pairs recorded by the subject. However, the determination of ($\underline{d}$) for any given spatial feature was complicated by the fact that the subject had complete freedom in the selection of the particular spatial features to be measured. It should be recalled that the subject was instructed to select the six strongest, i.e., most prominent, features. This instruction undoubtedly guided the subjects in their choice of features, particularly in the transform domain where the magnitude of each energy concentration was easily determined by the color coding. Nevertheless, it was necessary to establish a systematic procedure for matching each of the six pairs of coordinates selected by the subject with the most closely corresponding coordinates in the reference set before the ($\underline{d}$) values for the trial could be computed.

The automatic matching procedure was based on an algorithm which computed a matrix table of the ($\underline{d}$) values for all combinations of six reference coordinates against the six coordinates recorded by the subject. The algorithm then performed the matching by selecting the six lowest ($\underline{d}$) values in order, starting with the lowest value. As each match was made in this manner, the corresponding row and column of the matrix of ($\underline{d}$) values were eliminated prior to the selection of the next highest ($\underline{d}$) value. The result achieved by this procedure was the selection of a set of six ($\underline{d}$) values, the sum of which was lower than any other possible alternative matching. In other words, the subject's performance score for the trial was maximized in a consistent, objective manner by this procedure.

Additional features of the scoring algorithm included the inversion of the ($\underline{d}$) values to normalized units ranging from 0.0 (lowest score) to 1.0 (highest score) and the use of weighting factors applied to the converted ($\underline{d}$) values prior to averaging over the set of six. A more complete discussion of the scoring procedures may be found in appendix C.

SCORING CRITERION

As noted above, the spatial frequency coordinates selected by the subjects were scored against a set of reference coordinates. These reference coordinates were the best available estimates of the "true" coordinates for the six most prominent signal-related spatial features in each of the three image patterns. One of the difficulties inherent in the use of real acoustic data in studies of this nature is that "truth" with respect to signal characteristics is often difficult to define in any precise, quantitative way. Even in cases of high SNR level, the ambient noise in the ocean environment modifies the signal-related patterning in a random manner that cannot be specified in any exact mathematical sense.

The approach taken in this study to approximate the "true" coordinates was to derive them from the transform domain representation of the three image patterns at the highest SNR level, i.e., the image without added noise. Since the two-dimensional Fourier transform process decomposes an image in terms of its two-dimensional spatial frequency components, the transform provides a precise measure of the relative energy for all spatial frequencies in the original image. Through the use of color coding in the transform domain, spatial features in the original image can be individually identified and analyzed based on the location and magnitude of the spatial frequency energy concentrations.

There are other reasons why a more precise estimate of the "true" coordinates can be obtained from the transform rather than from any other feasible alternative (the original image, for example). The transform representation provides a display representing the decomposition of the spatial frequencies ordered by magnitude. In addition to ranking the features by prominence, the use of a color coding scheme designated "band-step coding" outlines the concentric contours for each peak, simplifying the localization of the coordinates of the highest point in the energy concentration (reference 2). Through the use of the interactive cursor with direct readout of the magnitude at the cursor location, an operator can easily determine the exact coordinates of the highest magnitude, even where the color coding provides no apparent differentiation. Finally, the operator is provided a direct readout of the coordinates for the cursor when he is satisfied that the peak has been located.

It should be apparent that there is no comparable method for systematically extracting and ranking spatial features using only information from the original image. The use of the interactive grid overlays does provide a direct readout of the spatial frequency coordinates of the overlay, but the spacing and slope of the overlay are determined solely by the operator's judgment of the best match to perceived pattern features.

For the cases of the highest SNR images, it was apparent from a preliminary examination of the recorded measurements that in the transform domain the coordinates of the six most prominent features were extracted with high consistency across the group of six subjects. For the comparable measurements in the image domain, a degree of consistency was obtained for the three most prominent spatial features only. In view of this high degree of consistency among the subjects in the transform domain, the decision was made to derive an estimate of the "true" coordinates by averaging the recorded measurements of the six subjects, obtained from the transform domain at the highest SNR level for each of the three image patterns. The resulting set of six coordinate pairs, ranked by the magnitude or prominence of the corresponding spatial feature, constituted the set of "reference"

---

[2] Naval Ocean Systems Center, NOSC TR 207, "Grey-Scale Versus Color Coding of Acoustic Data Images," by Robert S. French, March 1978.

coordinates used to score the six measurements made by each subject on each trial. These reference values were used for all the conditions tested under the experimental design, i.e., for both the image and transform modes and for all three SNR levels.

Use of the transform mode test results to derive the criterion for scoring imposes a constraint on the upper limit of the scores in the original image mode. The scores made by the subjects using the original image at the highest SNR level can be as high as, but no higher than, the scores made by using the transform. This constraint means that the results comparing the two modes should only be interpreted relative to each other, not as though the scoring criterion represented an absolute standard or "truth." Note that the scoring constraint imposed by the nature of the criterion applies only to the highest SNR level, as this level was used to derive the criterion. At the two lower SNR levels, there is no constraint on the possibility of higher performance scores in the original image mode.
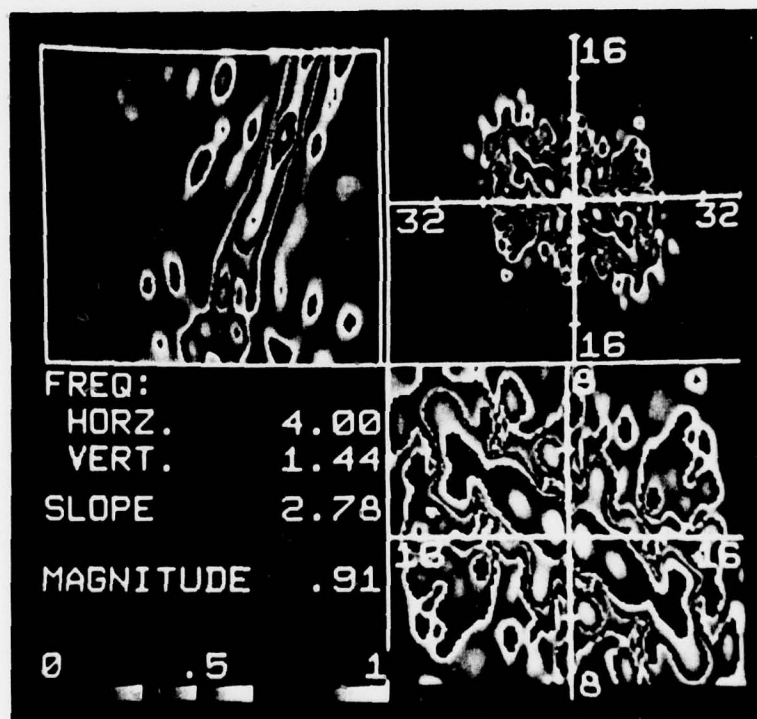
Figure 1. The combination format representing the transform mode of presentation. The original image is in the upper left quadrant; the transform, upper right; a 2-to-1 magnification of the transform, lower right; and a readout of data parameters, lower left.
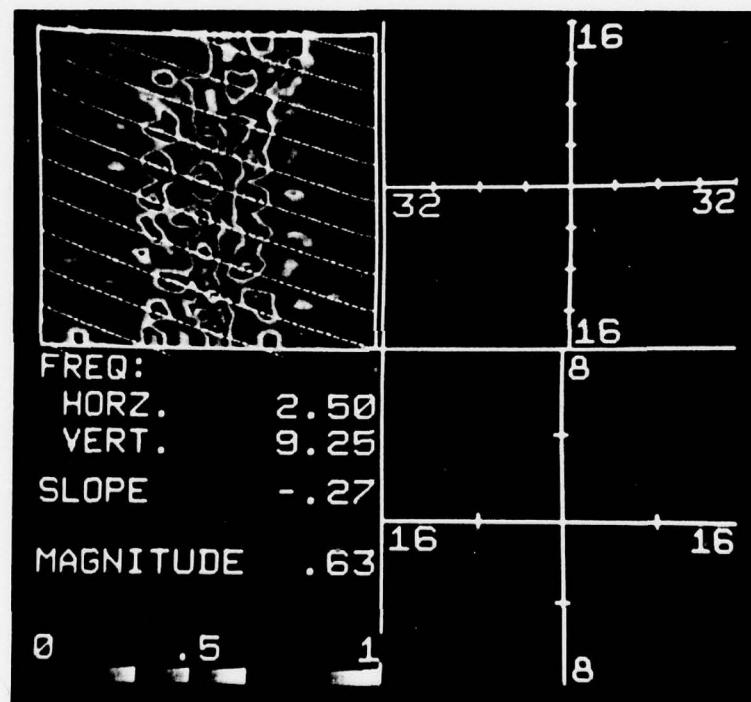
Figure 2. Format for the original image mode of presentation. An overlay of grid lines to measure the slope and frequency of a prominent periodic component is illustrated.
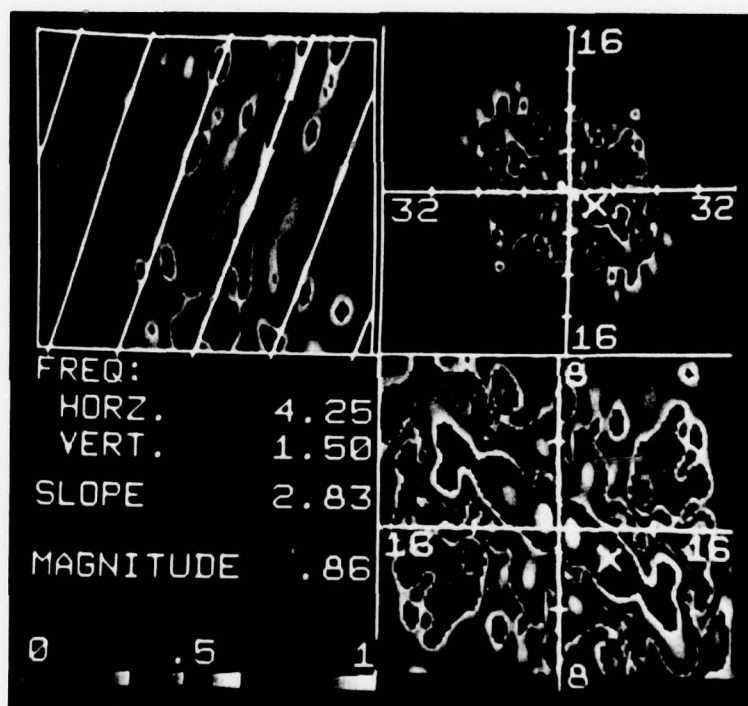
Figure 3. Format for the transform mode of presentation. A grid overlay coupled to a selected energy peak in the transform is illustrated. The slope and spacing of the overlay correspond to the vertical and horizontal spatial frequencies indicated by the "X" shown in the transform domain.
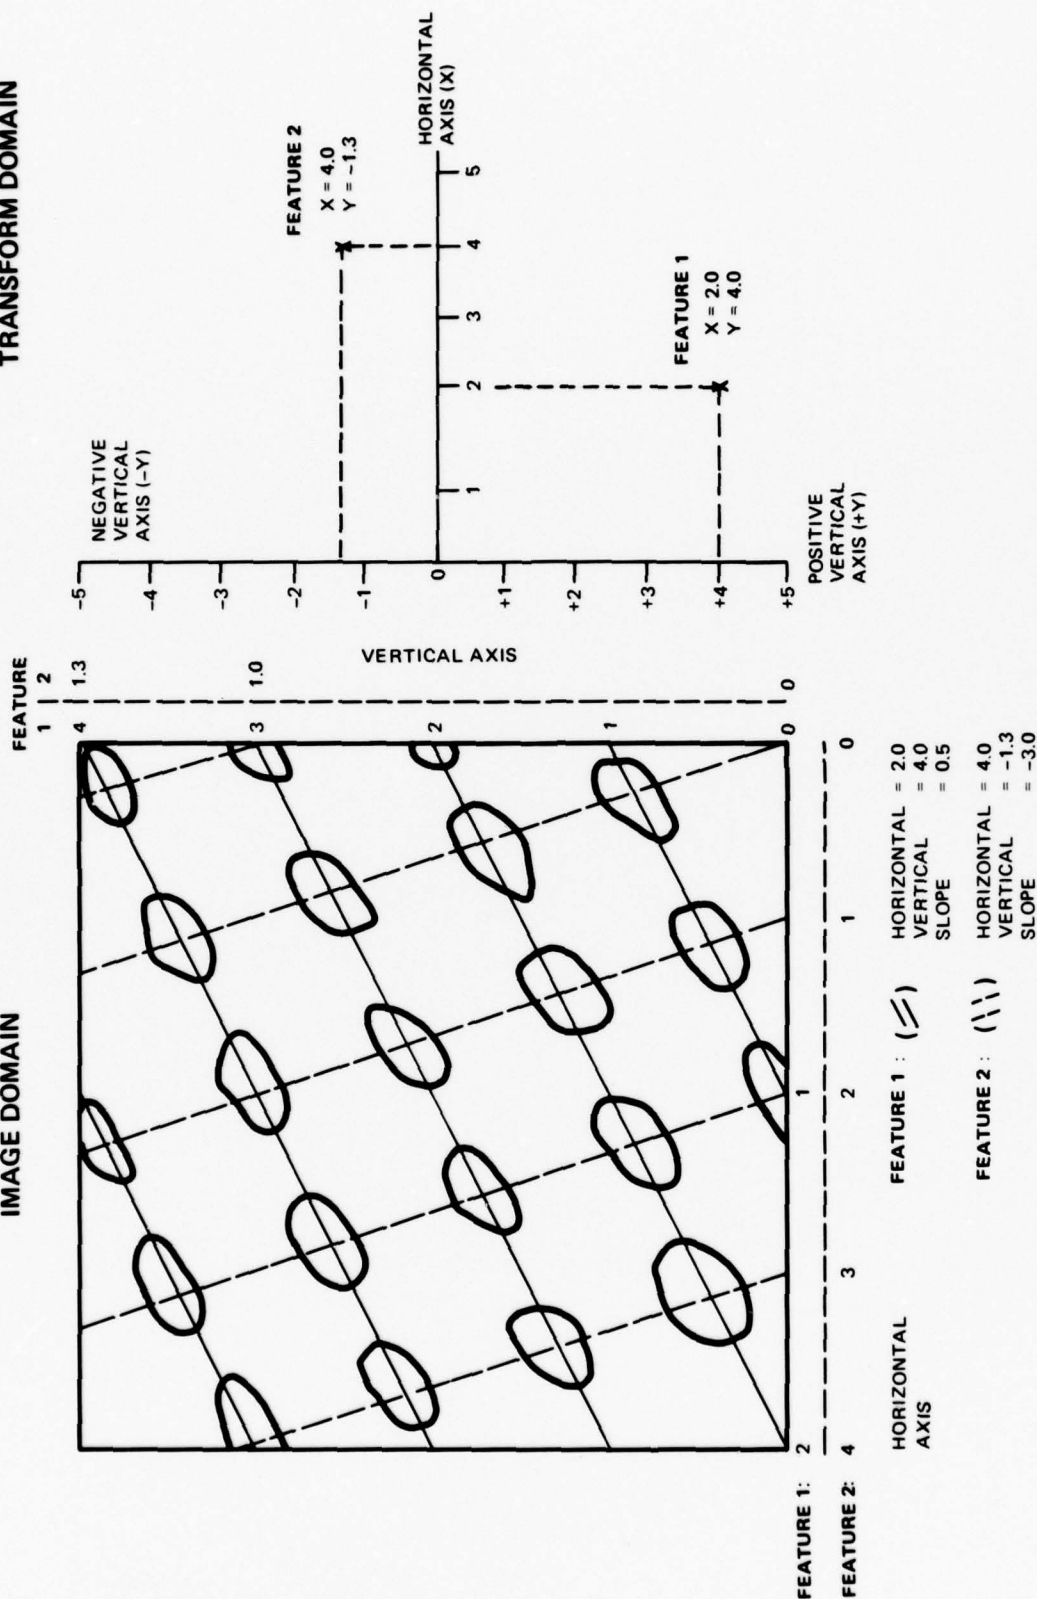
Figure 4. Graphic representation of two periodic spatial features with grid overlays as they might appear in the image and transform domains. The spatial frequency of each feature is defined in the image domain by the number of horizontal and vertical axis crossings (counting from the lower right corner of the image) and by the sign of the slope relative to the vertical axis. In the transform domain, the spatial frequency of each feature is represented by the location of the corresponding energy peak scaled in X,Y coordinates that correspond to the number of horizontal and vertical axis crossings. Because the transform is symmetric about the origin, only two of the four quadrants are shown in the transform domain.
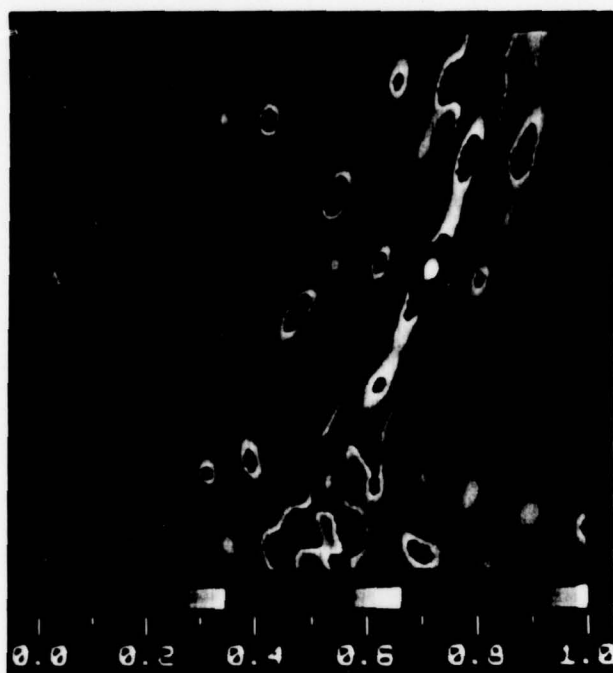
PART A. +5.8 dB
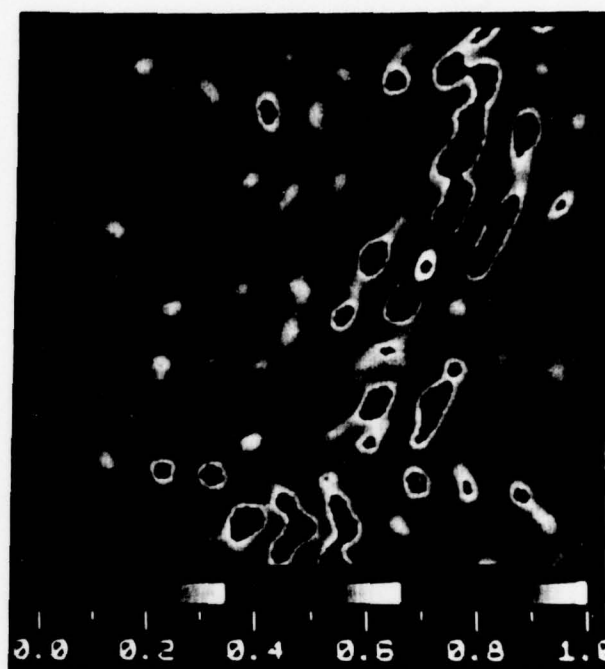
PART B. −5.1 dB
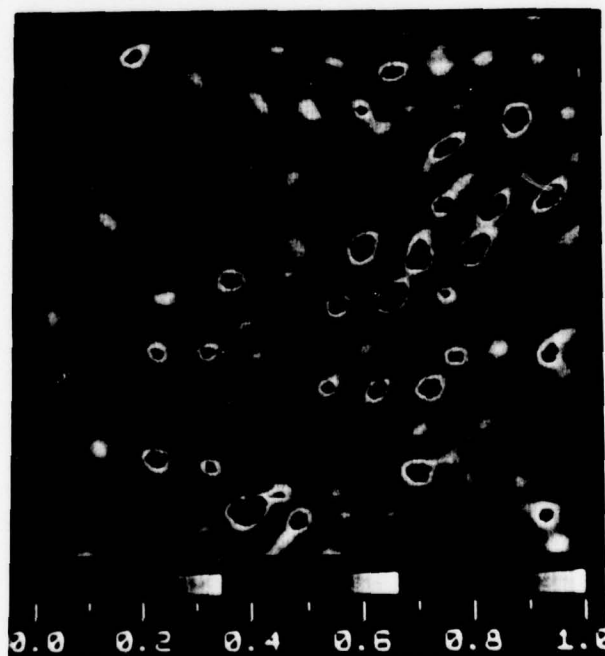
PART C. −9.3 dB

Figure 5. Test data set for pattern 1 at SNR levels of +5.8, −5.1, and −9.3 dB. Note the effect of the noise on the perceptibility of periodic spatial features that are readily apparent in the image before noise is added (5C).

21
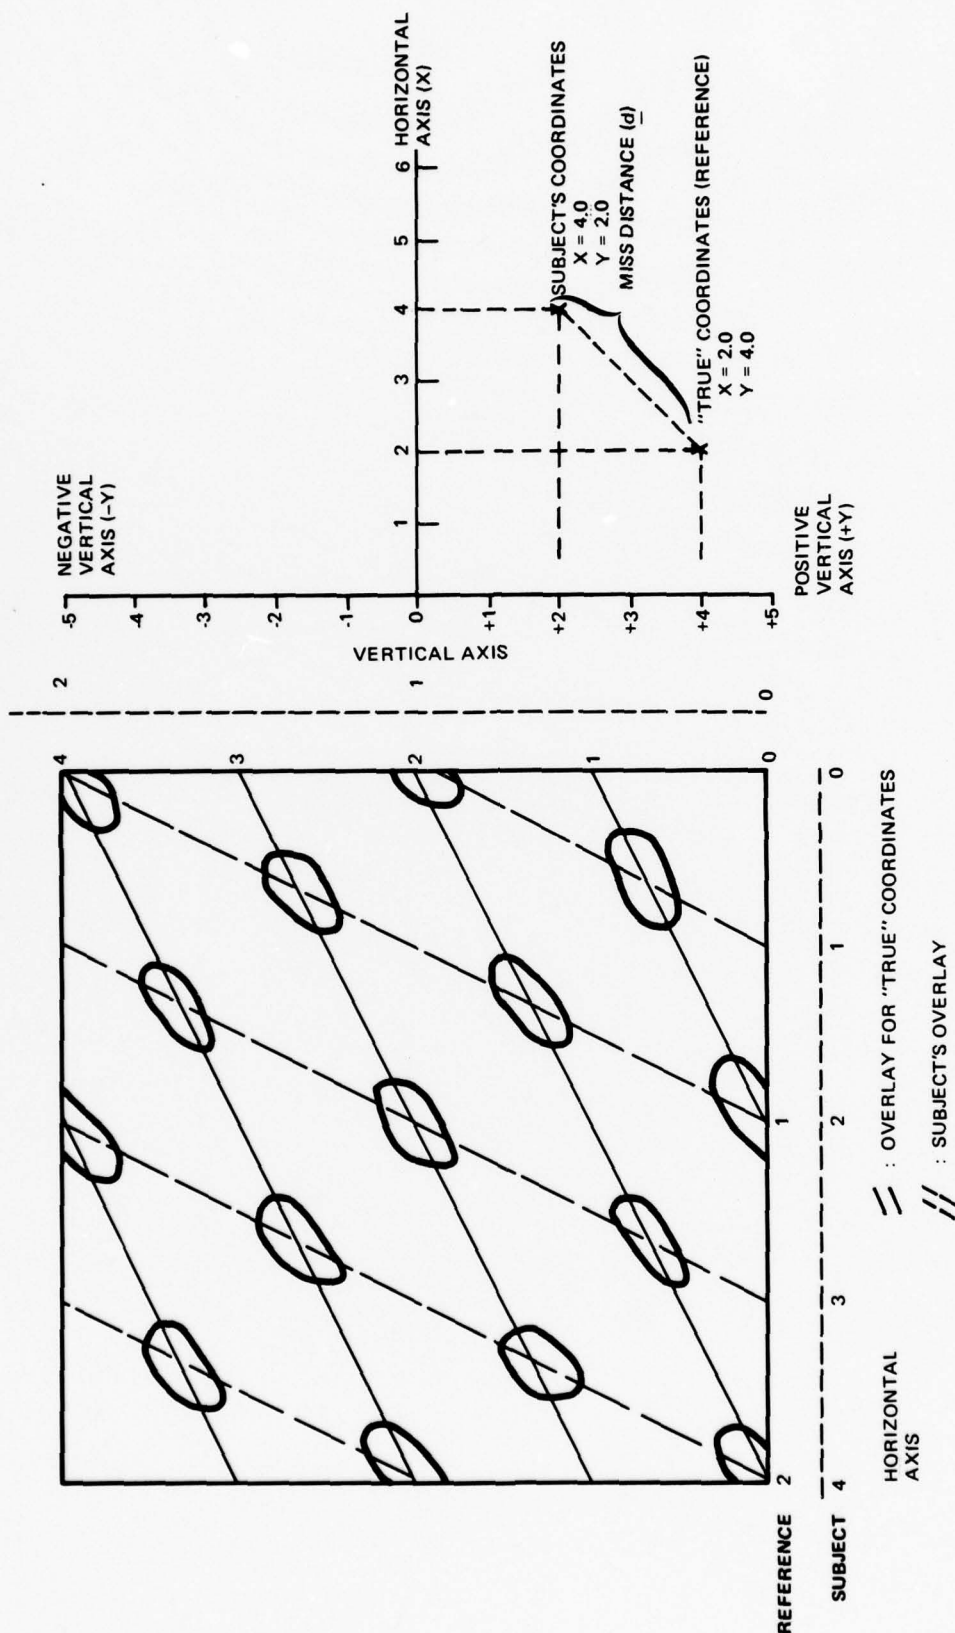
Figure 6. Graphic representation of the scoring concept of "miss" distance ($\underline{d}$). A grid overlay for an estimate of the "true" (reference) coordinates for a spatial feature and a hypothetical overlay constructed by a subject are shown in the original image domain. ($\underline{d}$) is defined as the absolute distance in the transform domain between the two points defined by the coordinates of the reference and the coordinates selected by the subject, respectively. The value of ($\underline{d}$) has been exaggerated for illustrative purposes. Because the transform is symmetric about the origin, only two of the four quadrants are shown in the transform domain.

23

RESULTS

The normalized mean performance scores for the six subjects comparing the two modes of presentation at the three SNR levels are in figure 7. The data for the three patterns have been combined in this figure. They indicate that at the highest SNR level (from +2.1 to +5.8 dB) the periodic components are measured with almost twice the efficiency when the transform is available (0.88 versus 0.50). However, when the SNR level falls to −5.1 dB the two performance curves converge and there is little difference in performance between the two modes. Finally, at the lowest SNR level, −9.3 dB, the original image alone yields slightly better scores than does the transform. The mean performance scores for the six subjects at the three SNR levels are in table 1. Figure 8 presents the same comparison between modes as shown in figure 7, but for each pattern separately. These data indicate that the difference favoring the use of the original image at the lowest SNR level is present in two out of three of the patterns, suggesting that the rate of convergence may be pattern related.

Table 1. Mean performance scores for the six subjects. The scores are tabulated to compare the original image and transform modes of presentation and the three SNR levels. Note that the average scores for the original image mode at -9.3 dB are slightly higher for each subject.

Value, dB								
−9.3			−5.1			+5.8, +2.1		
Subject	Original	Transform	Subject	Original	Transform	Subject	Original	Transform
1	0.24	0.17	1	0.38	0.39	1	0.51	0.86
2	0.28	0.18	2	0.44	0.38	2	0.45	0.86
3	0.22	0.21	3	0.29	0.39	3	0.58	0.91
4	0.23	0.18	4	0.31	0.42	4	0.52	0.88
5	0.38	0.15	5	0.43	0.39	5	0.53	0.87
6	0.27	0.18	6	0.27	0.40	6	0.44	0.88

To help understand the nature of the task confronting the subjects, two photographs for the transform mode of presentation for pattern 1 are in figure 9. (The comparable photograph for the highest SNR level is in figure 1.) It is evident from these photographs that as noise is added to the original image the number of energy concentrations in the transform increases. Since the order of presentation of SNR levels was from lowest to highest, i.e., the subject had not yet been exposed to the transforms with stronger signal-related peaks, the subjects had no basis for rejecting these noise peaks with high magnitude coding. Undoubtedly it is this proliferation of easily measured noise peaks which accounts for the relatively poor performance in the transform mode at the lower SNR levels.

A supplementary analysis was made to determine if the difference between the two modes was a function of the relative amplitude or prominence of the spatial feature. The scoring algorithm was modified to divide the six scores for each subject into two sets: those three scores matching the spatial features in the reference set with the three highest amplitudes and the remaining three which were the best match with the remaining low amplitude features. The results of the analysis are in figure 10. Each set of functions is now based on three scores for each subject instead of six. When comparing figures 10A and 10B for the

highest level it is evident that the average score in the original image mode is reduced by a greater amount (0.65 to 0.34) than it is in the transform mode (0.94 to 0.83), i.e., the less prominent periodic spatial features are particularly difficult to find in the original image as compared with the transform. For the two lower SNR levels there is essentially no difference between the two modes. The performance with both modes falls off dramatically for the less prominent spatial features when noise is added to the image.
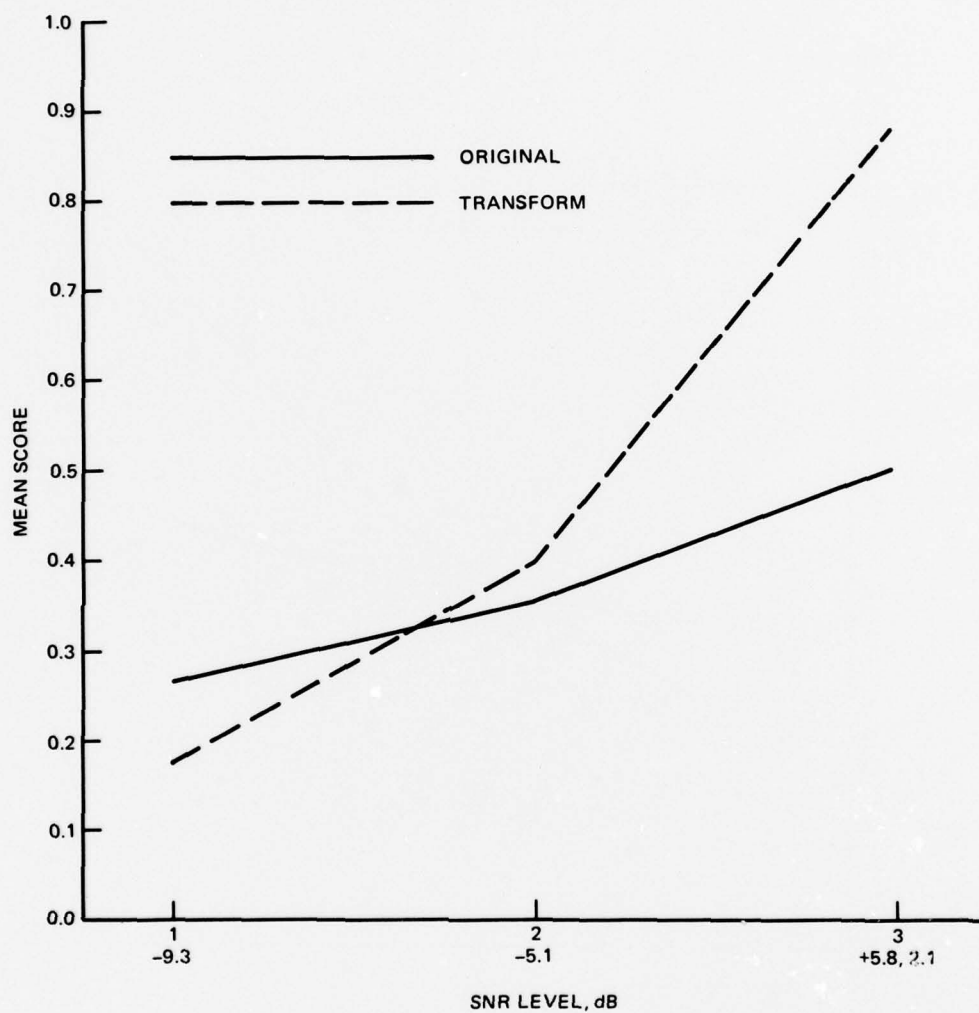
Figure 7. Normalized mean performance scores for six subjects comparing the original and transform modes of presentation at three SNR levels. The data for the three patterns are combined.
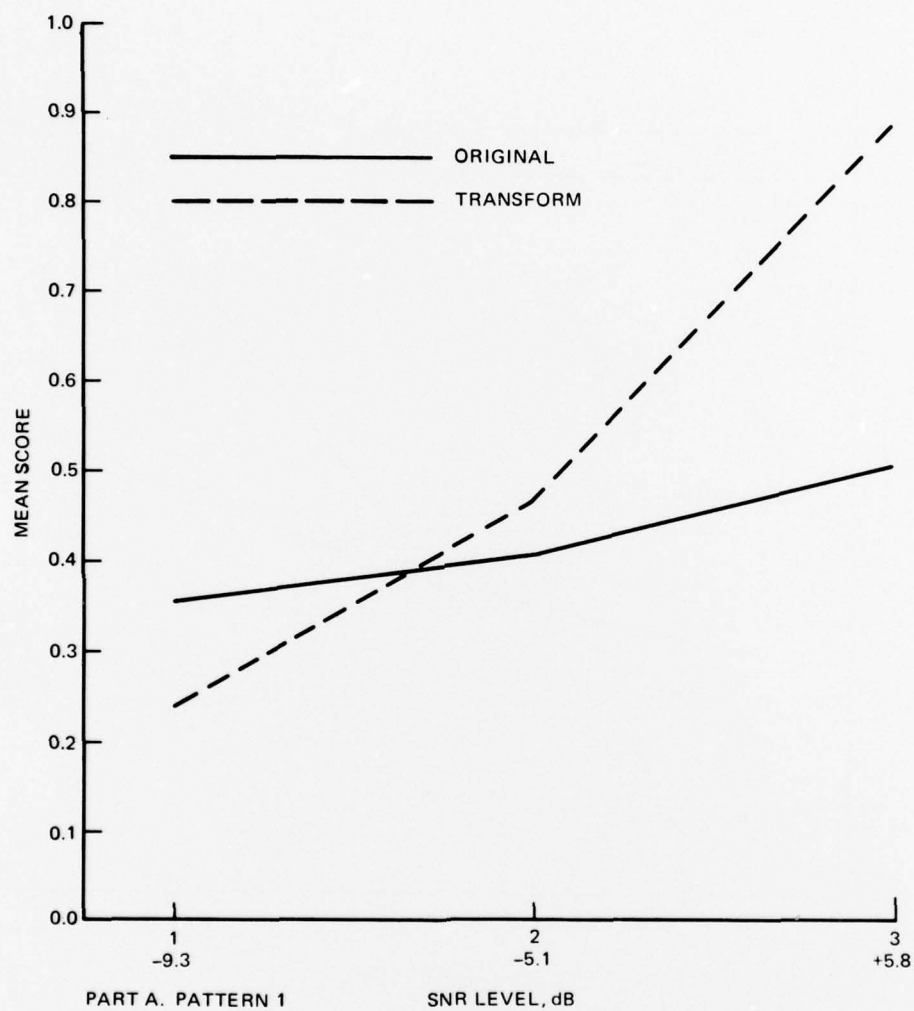
Figure 8. Normalized mean score for the six subjects comparing original and transform modes at three SNR levels.

28

Figure 8. Continued.

29

PART C. PATTERN 3

Figure 8. Continued.

30
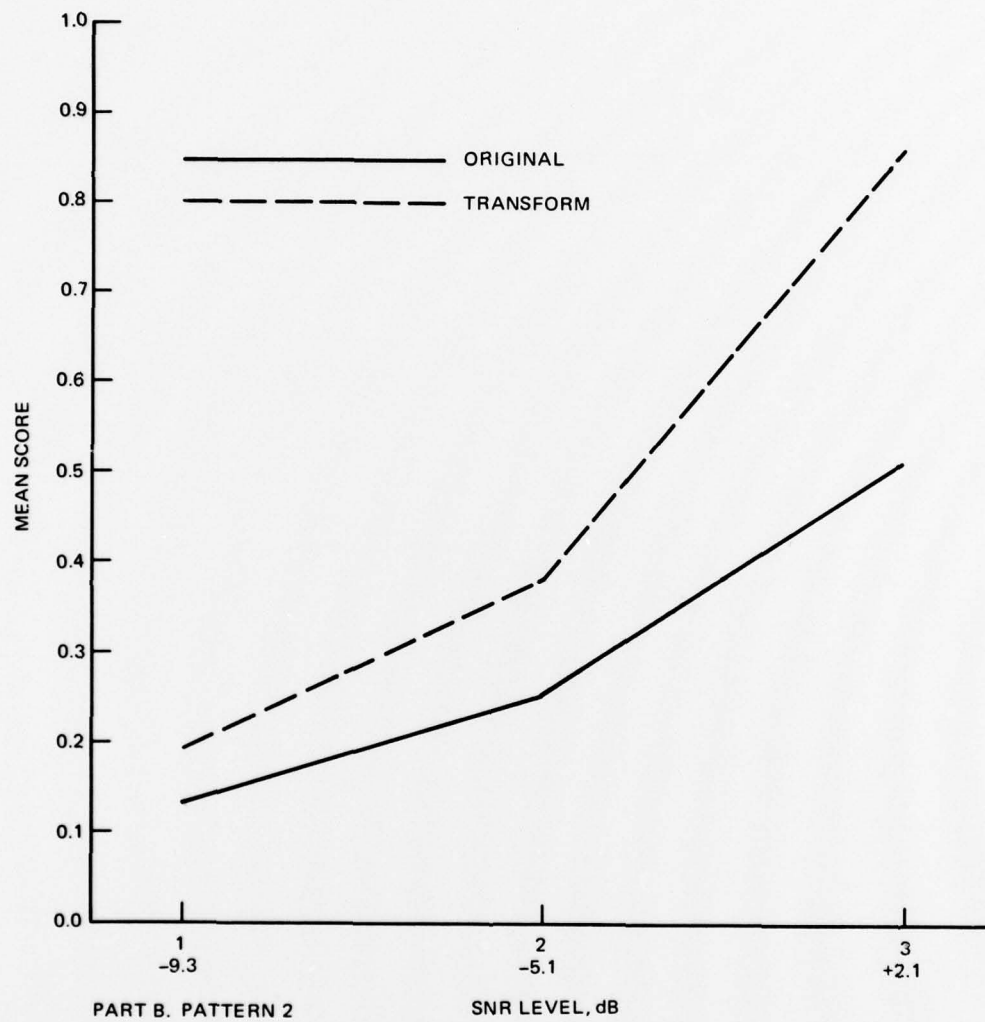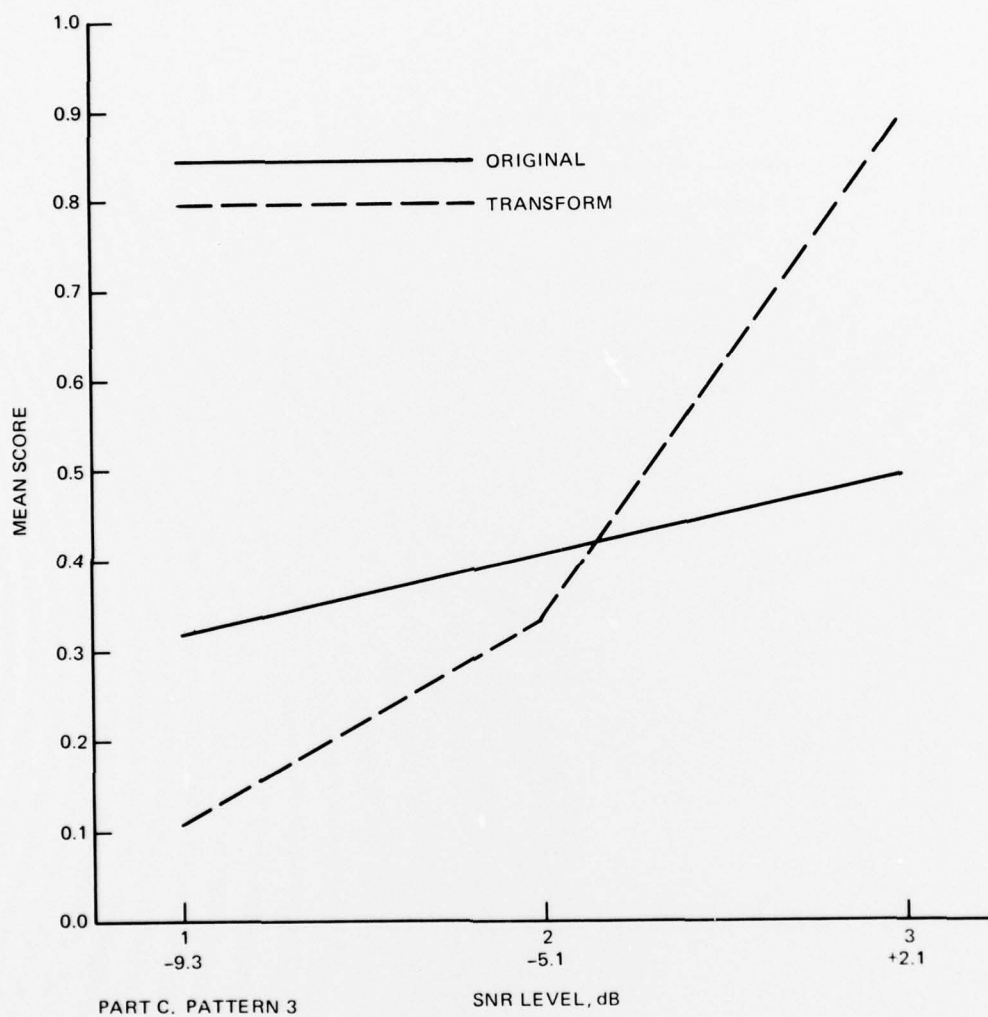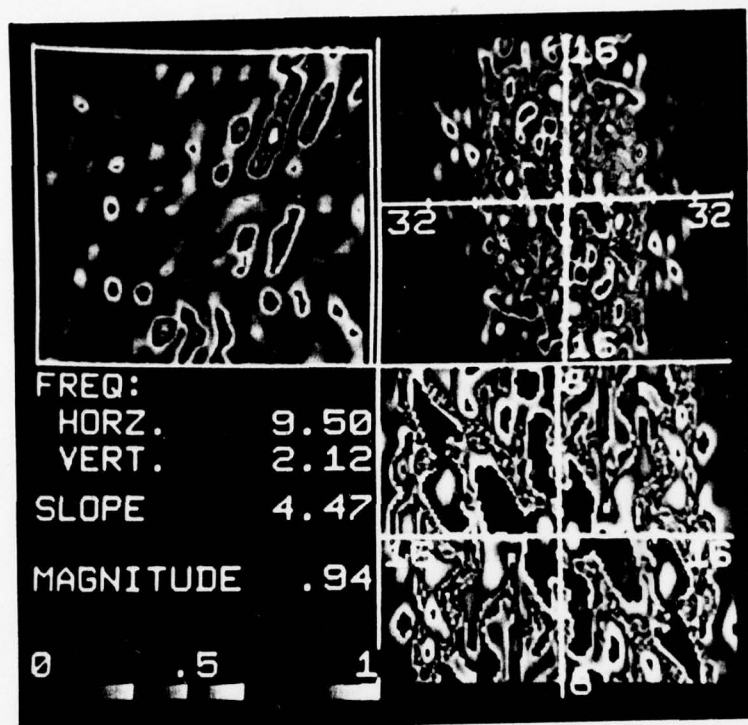
PART A. –5.1 dB
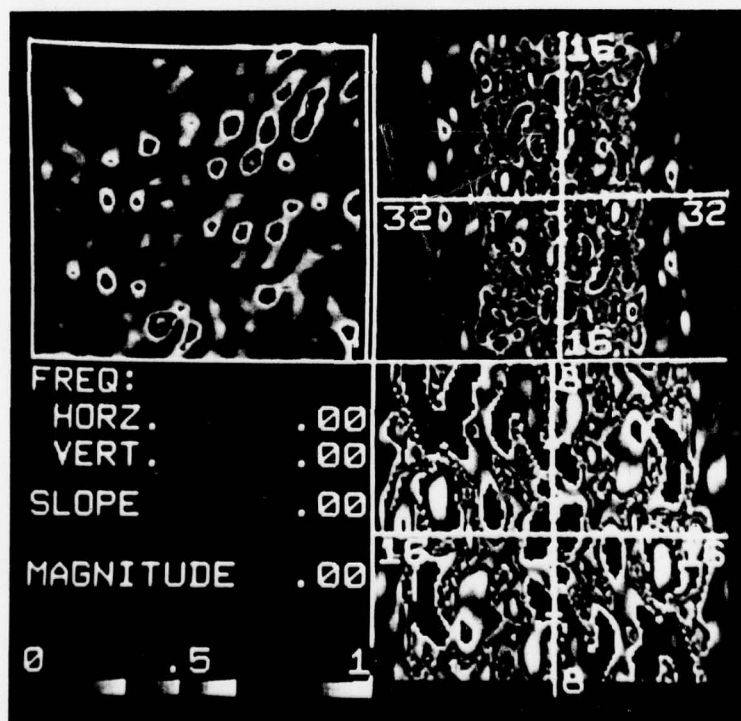
PART B. –9.3 dB

Figure 9. The transform mode of presentation for pattern 1 at the intermediate and lowest SNR levels. Note the greater number of noise peaks at the lower SNR level.
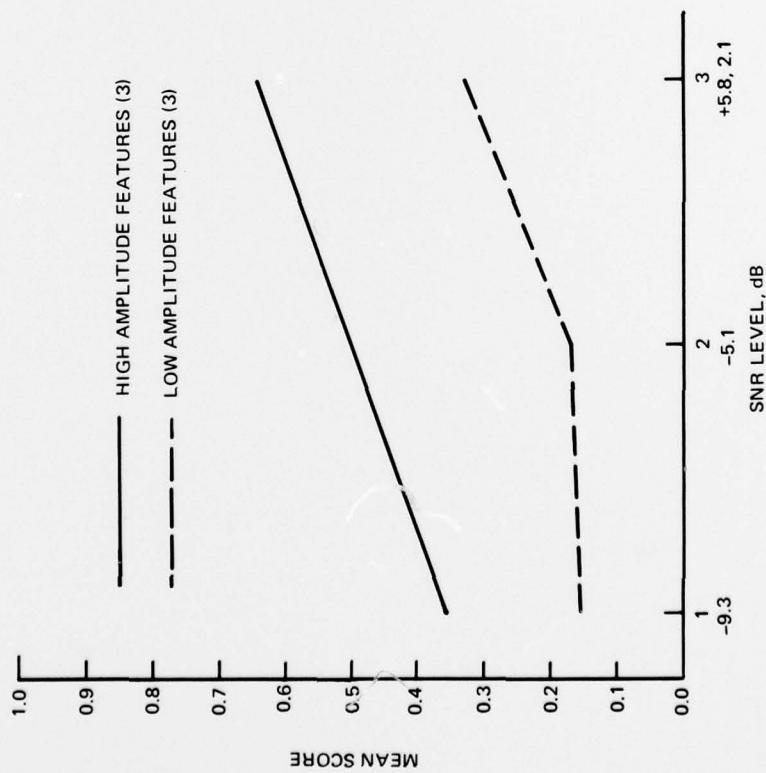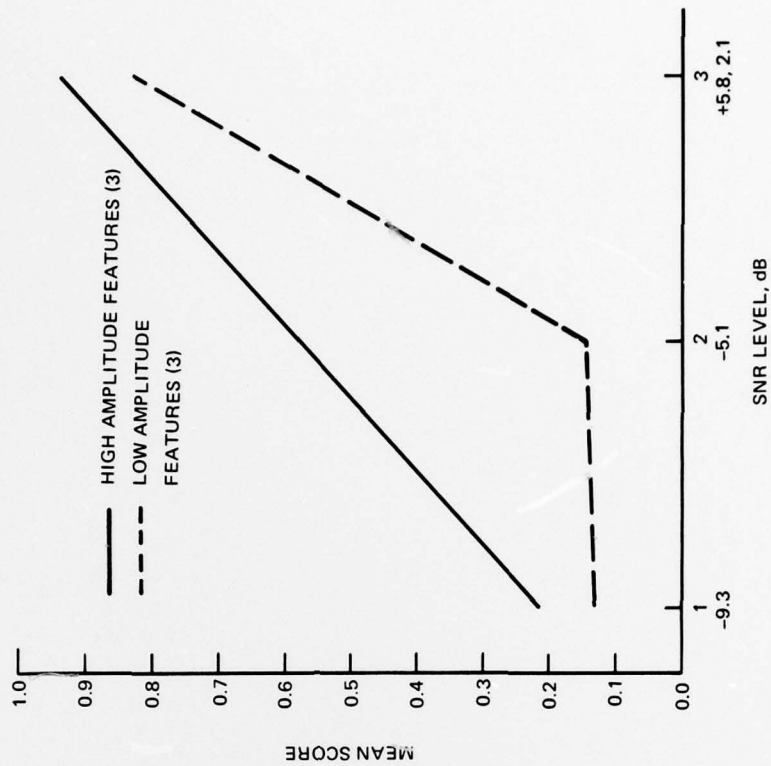
Figure 10. Modified normalized mean scores for the six subjects at three SNR levels. The six scores for each subject have been divided into two sets: the three highest amplitude (most prominent) spatial features and the three lowest amplitude (least prominent) spatial features.

33

CONCLUSIONS

At high SNR levels the transform domain image representation is clearly superior to the original image as a mode for systematic measurement and ranking of the periodic spatial features in the image. As the SNR level is reduced, the accuracy of measurement in the transform domain falls off more rapidly than in the original image domain, resulting in a convergence of the performance curves under the two modes. With the SNR level reduced so that signal-related patterning is barely perceptible, measurement of periodic spatial features in the transform domain is no better than in the original image domain. However, the amount of convergence for one of the patterns was less than for the others, suggesting that the rate of convergence may be pattern related.

A supplementary analysis was performed to identify factors that might account for the superior performance in the transform mode at the higher SNR levels. The magnitude or prominence of the spatial feature was found to be such a factor. At each SNR level, performance scores for the three most prominent spatial features were compared with the scores for the three less prominent features. At the highest SNR level, in the original image mode, the measurement accuracy for the less prominent features was significantly lower than for the three most prominent features. A comparable difference in accuracy was not found when measurement accuracies in the transform domain were compared: At the highest SNR level the less prominent spatial features were, relative to the most prominent features, more difficult to identify and measure in the original image domain than in the transform domain. At the two lower SNR levels no comparable difference between the two modes was observed. The measurement accuracy for the three less prominent periodic features was equally poor in both the original image and transform domain.

At high SNR levels the transform domain is an excellent mode for systematic measurement and ranking of both the major and the less prominent periodic spatial features using the interactive aids developed for this study. There exists no comparable method for systematically extracting and ranking periodic spatial features from the original image alone.

APPENDIX A. COMTAL DISPLAY SYSTEM

All experiments described in this report were conducted in the NOSC Surveillance Display and Image Processing Laboratory. The test data were viewed directly on the COMTAL 8300 Display System (figure A-1, center, back). This display system, which is the central feature of the laboratory, provides the user with the capability to display and manipulate interactively high-resolution color images. The host computer for the laboratory is the UNIVAC 1108. The COMTAL accepts input/output commands and data from the 1108 and internally performs scan conversion and refresh storage operations on the data to produce bright flicker-free displays in grey scale or color. Pertinent features of the COMTAL system include the following:

- 512- by 512-picture element (pixel) resolution with a 1:1 aspect ratio.
- Selectable data display modes including grey scale, pseudocolor, and true color.
- Three independent images stored at one time.
- Eight bits of intensity coding stored for each pixel (six bits displayed).
- Scan conversion for moving (rising or falling) raster presentations.
- Graphic overlay for superimposing outlines, grids, or alphanumerics on displayed images.
- Trackball positioning of a target pointer using a trackball input control device.
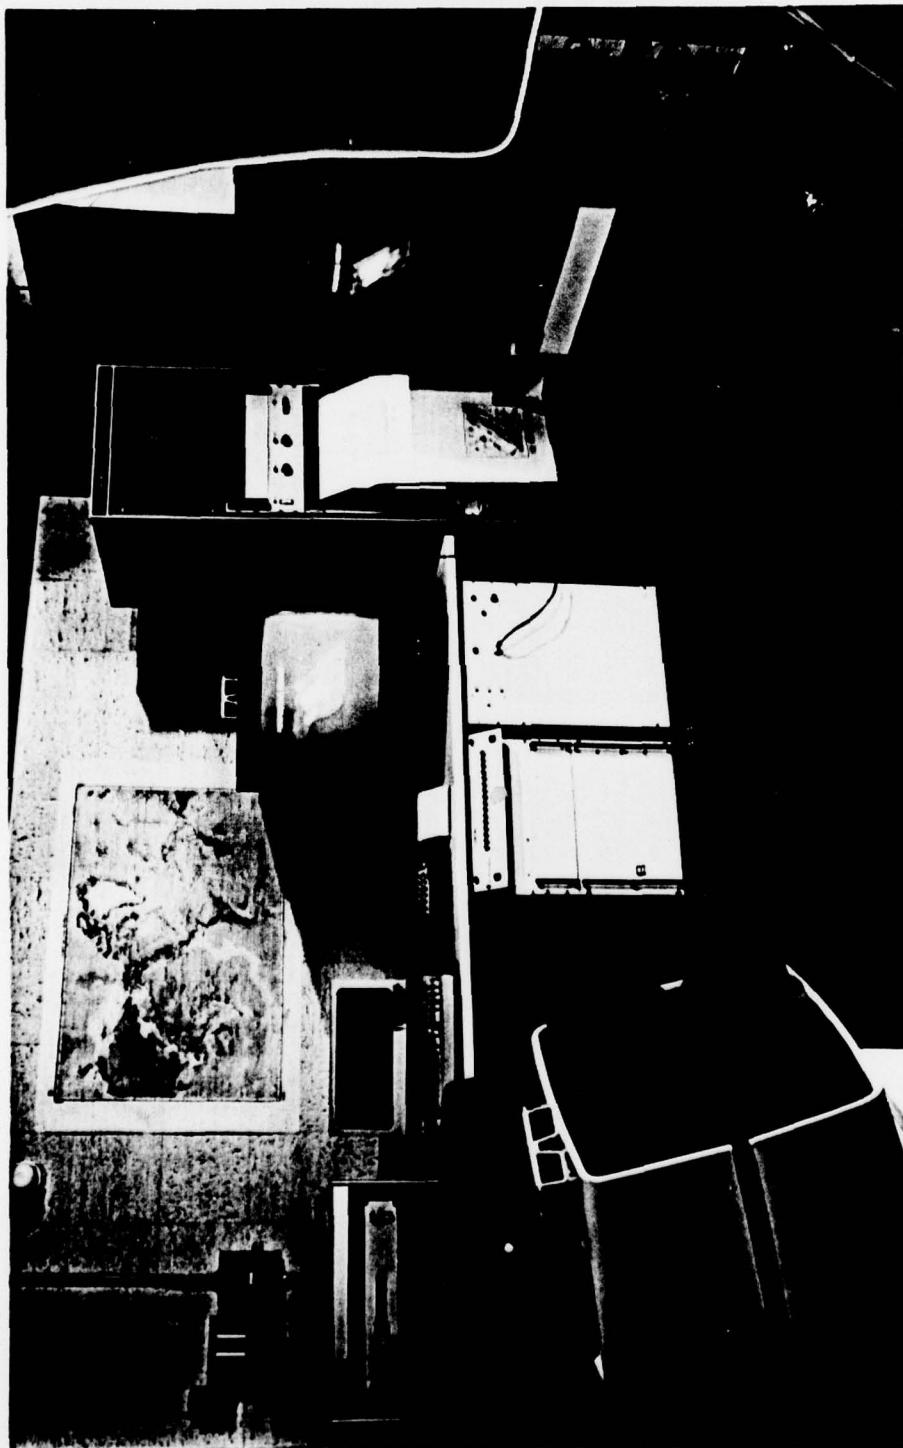
Figure A-1. NOSC Display and Image Processing Laboratory.

APPENDIX B. TEST INSTRUCTIONS

These are the exact instructions given to the six subjects at the beginning of the tests.

"This is an experiment to evaluate the two-dimensional Fourier image transform as an aid for measuring spatial frequency patterning. Much of this patterning is periodic in nature and gives rise to measurable concentrations of energy in the transform domain. These periodicities are often evident from direct observation of the original image, as well. In this experiment you will be concerned with measurement of the spatial frequency of axis crossings in the horizontal and vertical, the slope of the periodic pattern and in some cases the magnitude of the energy concentration in the transform. The definitions of each of these parameters are provided later in these instructions. The approach will be to compare the accuracy and completeness of the measurements under two conditions: a) directly from the original image, b) when viewing both the original image and the transform. The images will be viewed on the COMTAL in a standard format. The original image is in the upper left quadrant; the two-dimensional Fourier transform in the upper right quadrant; a 2:1 magnification of the transform is in the lower right quadrant; and the spatial feature parameters are summarized in the lower left quadrant. Because the transform is symmetric with respect to the original, measurements may all be made in the two subquadrants *within* the transform area to the right of the axis crossing (zero spatial frequency or "DC"). The major features of this formature are illustrated in Figure 1." (Same as figure 1 in this report.)

"Various readouts and overlays will be available to aid you in measuring the pattern parameters. When you are working in the image domain, your primary measurement aid will be to generate overlays by positioning the CURSOR. For each overlay the computer will readout the spatial frequency components in the horizontal and vertical (i.e., the number of lines crossing the horizontal and vertical axes) and the slope of the overlay." (See figure 2 in this report.)

"When both the original image and transform are viewed, a variety of CURSORS and readouts are available which directly couple the two image domains. For example, when you generate an overlay on the original image, you will automatically generate an "X" in the transform at the coordinates defined by the overlay. You will also get a readout of the horizontal and vertical frequency components, the slope and the magnitude. When you select any spot in the transform using the CURSOR, the computer will generate the corresponding overlay in the original image as well as provide the readouts. Because the phase information is not used in the display of the transform, the overlay is not likely to be superimposed on the designated periodic feature. As an aid in verifying that you have located the correct energy concentration in the transform, you may reposition the overlay by placing the CURSOR at any point through which you wish an overlay to pass. Repositioning the overlay is done by shifting the entire overlay to the right or left. Neither the slope nor the vertical placement is changed by this action." (See figure 3 in this report.)

EXPERIMENTAL DESIGN

"The test sequence will consist of measurements on three different patterns using the two modes of image presentation, at three levels of signal strength. The experimental design is 3 (patterns) X 2 (modes) X 3 (SNR levels) = 18 trials. The test sequence will be the same for each subject and can be defined in terms of the rapidity with which the three variables are changed. The most rapidly changing variable will be the pattern, then the two modes and finally the SNR. The two modes will be presented in the order: original image, both. The three SNR levels will be presented in the order from lowest to highest.

"Your task in this experiment will be to determine and record on the prepared data sheets as many periodic spatial frequency components as you can find (up to six) in either the original image, or when both are viewed. Your measurements will be scored both on their accuracy relative to the correct values, and the extent to which you identify *all* six of the components. *Be sure* to include the six with the *strongest* patterning in the original image or the highest peaks in the transform. Do *not* include the DC component which is located at the origin.

"Since you will be viewing the same patterns repeatedly, under different conditions, it is important to minimize carryover from one set of measurements to the next. To minimize carryover you will not be allowed to compare measurements for a given pattern under the two modes or SNR conditions. Once the set of measurements for the three patterns has been recorded you cannot change your measurements or refer to them again. Each completed data sheet must be turned upside down and must not be looked at again during the test."

IMAGE DOMAIN MEASUREMENT

"To construct an overlay in the original image, you need to define three points using the CURSOR. Initially, the CURSOR is located at the lower left hand corner of the image. To construct an overlay, place the CURSOR over a prominent peak which you wish to have coincide with a line of the overly and press the SPACE BAR. The instruction will appear on the terminal to select another point on the same line. To do this, reposition the CURSOR to a second peak which together define the desired line and press the SPACE BAR, again. The next instruction on the terminal will be to select a position on a different line. The third point should coincide with a parallel set of peaks or pattern feature, with a spacing representing the interpeak distance and thus the spatial frequency of the desired overlay. Pressing the SPACE BAR the third time will lead to the construction of the overlay and will also give a readout of the horizontal and vertical spatial frequency components and the slope of the lines. If you wish to erase the overlay before constructing another one, press the "RUB OUT" button."

TRANSFORM DOMAIN MEASUREMENT

"To read out the parameters of any point in the transform, place the CURSOR over the peak of interest and press the SPACE BAR. You will get a readout of the horizontal and vertical coordinates of the point, the magnitude, and the slope of the corresponding overlay.

"Note that the slope of any point in the upper right quadrant of the transform is negative, increasing in value as you approach the horizontal axis. Points in this quadrant correspond to overlays in the original image that are inclined to the left of the vertical. In a similar manner points in the lower right quadrant are positive and correspond to overlays inclined to the right of the vertical. This change of slope "sign" occurs from original image to transform because the effect of the transform is to rotate the pattern 90 degrees. It should be noted that slope is always defined relative to the overlay in the image domain. An overlay consisting of horizontal lines has a zero slope (Y=inf. X=0.0); an overlay consisting of nearly vertical lines will have a very large value of slope. If the lines are fully vertical the slope is infinitely large but will be computed as zero, due to a computer limitation."

COMBINED IMAGE/TRANSFORM MEASUREMENT

"To construct an overlay and readout parameters when both the original image and transform are viewed, one would normally begin by placing the CURSOR over a prominent peak in the transform and pressing the SPACE BAR. This action will generate an overlay in the original image and give a readout of x, y slope, and the relative magnitude of the energy concentration. To reposition the overlay, place the CURSOR over a point in the original image through which you wish the overlay to pass and press the "LINEFEED" button. If you wish to remove the overlay to construct another one, press the "RUB OUT" button. *You may also begin* the analysis by constructing an overlay in the *original image* as in the image mode. Now, however, the image overlay will be coupled with the construction of an "X" at the corresponding coordinates in the transform. This coupling will allow you to verify that your overlay corresponds to an energy peak, and if not, to modify the overlay until it does correspond."

CONTROLS

"As a convenience, the use of each of the spatial keyboard entries and controls is summarized below:

RETURN	—	Sequences to the next trial
TRACKBALL	—	Controls the position of the cursor
SPACE BAR	—	Enters a point to define an overlay
LINEFEED	—	Repositions the overlay
RUB-OUT	—	Erases the overlay"

APPENDIX C. SCORING ALGORITHM

The subjects' recorded measurements were scored in terms of accuracy and completeness relative to a set of reference coordinates derived by measurement of the six strongest periodic spatial frequency components in the images. The coordinates for the set of reference values were obtained from measurements in the transform domain at the highest SNR level for the three patterns. The measurements made by the six subjects for these images were averaged to provide the best available estimate of the "true" coordinates of the six strongest components.

The first element in the scoring algorithm was the computation of the absolute "miss" distance ($\underline{d}$) relative to the reference coordinates for each of the six frequencies recorded by the subject. This aspect of the algorithm is complicated by the need to make the best possible match between the subject's six values and the reference set. The approach taken calculated all ($\underline{d}$) values for a matrix of the six reference values against the values recorded by the subject. The algorithm then performed the matching by selecting the six lowest ($\underline{d}$) values in order, starting with the lowest value. As each match was determined, the corresponding row and column of the matrix of ($\underline{d}$) values were eliminated from further consideration. The subject's performance score was maximized in a consistent objective manner by this procedure.

Following the computation of the six ($\underline{d}$) values for each trial, a cutoff value was determined to ensure that the data points included in the score were within a reasonable distance of the corresponding reference points. Based on a preliminary examination of the test data, it was determined that a cutoff value of 3 ($\underline{d}$) units would adequately separate near misses from cases of obvious mismatch. Accordingly, the next computational step was to subtract each of the selected ($\underline{d}$) values from 3. The resulting values were then divided by 3 to produce a normalized miss distance, Q_n, with units ranging from 0 (any ($\underline{d}$) value which is 3 or higher) to 1.0 (perfect score).

The final step was to sum the six values of Q_n. Prior to summing, each Q_n was multiplied by a weighting factor, W_n, proportional to the amplitude of the transform peak for the corresponding reference value. The weighting factor was determined by summing the amplitude values of the six reference values and dividing each amplitude by this sum. Since the weighting factors equal unity, the sum of the six scores when multiplied by the weighting factors also ranges between 0.0 and 1.0.

The mathematical relationship for computing the score for subject, S_k, on a given trial was

$$\text{score } (S_k) = \sum_{n=1}^{6} Q_n W_n \quad , \tag{C-1}$$

where

$\quad$ n = index for the selected coordinates

$$Q_n = \begin{cases} 1 - \dfrac{d_n}{3}, & \text{if } 0 \leqslant d_n \leqslant 3 \\[2ex] 0 & \text{if } d_n \geqslant 3 \end{cases}$$

$$d_n = \left[\left(X_{\text{reference}_n} - X_{\text{subject}_{k_n}} \right)^2 + \left(Y_{\text{reference}_n} - Y_{\text{subject}_{k_n}} \right)^2 \right]^{1/2}$$

and

$$W_n = \frac{a_n}{\displaystyle\sum_{r=1}^{6} a_r} \quad ,$$

where

   a_n = peak amplitude for a given reference value.

In equation (C-1) Q_n is the normalized miss distance, d_n is the absolute miss distance between the subject's test coordinates ($X_{\text{subject}_{k_n}}$, $Y_{\text{subject}_{k_n}}$) and the reference coordinates ($X_{\text{reference}_n}$, $Y_{\text{reference}_n}$), and W_n is a fractional weighting factor relating the peak amplitude for a given reference value (a_n) to the sum of the amplitudes of the six reference values, where r indexes the six reference values.

APPENDIX D. SUBJECTS

The type of experimental design used in this study ensures that differences between the subjects in background and experience will not affect comparisons involving the primary variables. However, it is desirable that the subjects possess skills and knowledge similar to those of the operational personnel towards which the study conclusions are directed.

The six subjects used in the experiment were drawn from the professional staff of the NOSC Signal Processing and Display Division. All subjects had served in a previous experiment comparing grey-scale and color coding of acoustic data images.[D1] Because the subjects had had previous exposure to the three images they were generally familiar with the patterning and the degraded appearance of the images at lower SNR levels. The subjects had a basic understanding of the Fourier transform as a result of their signal processing background, but had had only incidental exposure to any specific coupling of grid overlays to peaks in the transform domain for the test images used. Although the subjects as a group were highly trained in the theoretical aspects of signal processing, this knowledge was not required for the successful execution of the task. With respect to the skills and knowledge essential to the task, the subjects were comparable to operational Navy personnel who might utilize similar displays for data analysis at some future time.

[D1] Naval Ocean Systems Center, NOSC TR 207, "Grey-Scale Versus Color Coding of Acoustic Data Images," by Robert S. French, March 1978.